arXiv:1704.05001v1 [math.NA] 17 Apr 2017

REDUCTION-BASED ALGEBRAIC MULTIGRID FOR UPWIND DISCRETIZATIONS*

THOMAS A. MANTEUFFEL †, STEPHEN MCCORMICK †, STEFFEN MÜNZENMAIER ‡, JOHN RUGE †, AND BEN S. SOUTHWORTH §

Abstract. Algebraic muligrid (AMG) is a go-to solver for symmetric positive definite linear systems resulting from the discretization of elliptic PDEs, or the spatial discretization of parabolic PDEs. For diffusion-like problems, the time to solution typically scales linearly with the number of unknowns. However, convergence theory and most variations of AMG rely on A being symmetric positive definite. Hyperbolic PDEs, which arise often in large-scale scientific simulations, remain a challenge for AMG, as well as other fast linear solvers, in part because the resulting linear systems are often highly nonsymmetric. Here, a new reduction-based AMG method is developed for upwind discretizations of hyperbolic PDEs and other matrices with triangular structure, based on the concept of "ideal restriction" (AMGir). Theory is presented motivating the reduction-based approach and connecting it in a big-picture sense to other nonsymmetric AMG methods. The model problems considered are variations of a steady-state transport equation, with source-terms, discontinuities, and non-constant flow. AMGir is shown to be an effective and scalable solver for various discontinuous, upwind discretizations, with unstructured meshes, and up to 6th-order finite elements, expanding the applicability of AMG to a new class of problems. AMGir is implemented in PyAMG [8] and available at https://github.com/ben-s-southworth/pyamg/tree/air.

1. Introduction. Algebraic multigrid (AMG) is among the fastest class of algorithms for solving linear systems that result from the discretization of partial differential equations (PDEs) of elliptic type [10]. When applicable, AMG converges in linear complexity with the number of degrees-of-freedom (DOFs), and scales in parallel like $O(\log_2(P))$, up to hundreds of thousands of processors, P [7]. Originally, AMG was designed for symmetric positive definite (SPD) linear systems, and performs best when applied to the discretization of elliptic PDEs, or the spatial discretization of a parabolic PDE in time. Although this constitutes a large class of problems that arise in scientific simulations, many problems of interest remain that are difficult to solve for AMG, or any other "fast" solver. Because of this, solutions to such problems are generally obtained through Krylov methods or direct solves, each of which are more computationally expensive than desired. Efforts have been made to develop more robust AMG solvers applicable to difficult SPD systems such as anisotropic diffusion [33, 51], and nonsymmetric matrix equations [13, 33, 38, 50, 56], among others. Despite improvements in the state-of-the-art, one class of problems that comes up in scientific computation, but lacks robust algebraic solvers, are upwind discretizations of hyperbolic PDEs.

In contrast to elliptic and parabolic PDEs, the solution of hyperbolic PDEs lies on characteristic curves, and the solution at any point depends only on the solution upwind along its given characteristic. This allows for very steep gradients or "fronts" to propagate through the domain in the direction of characteristics. Typical continuous finite-element or finite-difference discretizations often struggle to capture such behavior because fronts are effectively discontinuities in the discrete setting, and not neces-

*This work was performed under the auspices of the U.S. Department of Defense through the National Defense Science & Engineering Graduate Fellowship Program, the U.S. Department of Energy under grant numbers (SC) DE-FC02-03ER25574 and (NNSA) DE-NA0002376, and Lawrence Livermore National Laboratory under contract B614452.

†Department of Applied Mathematics, University of Colorado at Boulder.

‡Fakultät für Mathematik, Universität Duisburg-Essen.

§Department of Applied Mathematics, University of Colorado at Boulder (ben.s.southworth@gmail.com).

sarily grid-aligned. It is worth noting that recent work introducing a functional scaling for a continuous least-squares finite-element discretization suggests that continuous discretizations can effectively capture fronts in the steady-state transport equation [32]. Nevertheless, due to the discontinuous-like behavior and the flow of information in a single direction along characteristics, discontinuous upwind discretizations are a natural approach to discretizing many hyperbolic PDEs [14, 35, 36, 45]. For a fully upwinded discretization, the resulting matrix has a block-triangular structure in some (although potentially unknown) ordering. Although solving a block-triangular matrix is an easy task in the serial setting, current approaches in parallel are limited by scalability.

Solving triangular systems in a traditional context using a forward or backward solve is a strictly sequential operation and, thus, does not scale well in parallel. In the dense matrix setting, this offers limited opportunity for parallelism. However, for sparse matrices, scheduling algorithms have been developed that can add some level parallelism to this process [3, 5, 28, 29]. In a sparse triangular matrix, there exists at least one DOF, say i_0, that is independent of other DOFs and can be eliminated from the system. Then, by nature of a sparse matrix, we expect multiple DOFs to depend only on i_0, and can then be computed in parallel. In a discretization sense, these would be DOFs immediately downwind from i_0. Once new DOFs have been computed, DOFs downwind of them can be computed, and so on until the solution is obtained. This algorithm requires only $2 \cdot nnz$ floating point operations for a sparse matrix with nnz nonzero entries, but the parallelism depends on the structure of the matrix. Much of the research on increased parallel performance for such problems is focused on shared memory environments such as GPUs [28, 29].

One problem and discretization that requires many parallel sparse triangular solves is the so-called "diffusion synthetic acceleration" (DSA) approach to solving the linear Boltzmann transport equation (so-called "synthetic acceleration" is simply preconditioning) [1, 16, 26, 45]. The full linear problem is seven dimensional, necessitating highly parallel solvers for even moderately refined grids, and is fundamental to studying neutron and radiation transport. DSA is effectively a two-level multigrid algorithm. A "transport sweep" acts as a surrogate for relaxation, and consists of discretizing in angle and solving a steady-state transport equation over all angles in the spatial domain. An accurate transport sweep puts error in the range of a surrogate "coarse-grid" diffusion solve, and the process iterates back and forth [1]. AMG is a fast and scalable solver for the diffusion solve, but the transport sweep consists of solving many discretizations of a hyperbolic PDE, often based on an upwind discretization [14, 35, 36, 45], and is difficult for most iterative solvers [52].

Similar scheduling algorithms to those used for sparse triangular solves on GPUs have been developed in distributed memory environments to perform "transport sweeps" [2, 6, 16, 26]. On a structured mesh, these algorithms were shown to have an optimal theoretical scaling over all angles of $O(dP^{\frac{1}{d}})$, for P processors and a problem of dimension d [2]. Although this is reasonable scaling in parallel, it is suboptimal on the so-called "road to exascale," as a steadily increasing number of processors become available. On irregular meshes, scheduling for the transport sweep becomes more difficult and this theoretical bound cannot be achieved. The sweeping algorithm has been shown to be a bottleneck in large simulation codes [44]. On the other hand, AMG scales like $O(\log_2(P))$ in parallel [7] and $O(n)$ with respect to DOFs, even on irregular meshes. An effective AMG algorithm for triangular systems could prove faster than known scheduling algorithms for transport sweeps in a highly parallel setting. For

example, for $P = 10^6$ and $d = 3$, we have $dP^{\frac{1}{d}} = 300$, while $\log_2(P) \approx 20$, yielding a factor of 15 improvement in parallel efficiency. Of course the leading constants may change this balance in either direction, especially on irregular meshes, but this suggests there is a reasonable P for which AMG will outperform a traditional transport sweep.

Traditional AMG is based on minimization principles in the A-norm or *energy-norm*, $\|\mathbf{x}\|_A^2 = \langle A\mathbf{x}, \mathbf{x}\rangle$. If A is not SPD, such a norm is not well-defined, and the bulk of theory motivating AMG (for example, [10, 22, 55]) no longer holds. Nevertheless, AMG has still been applied to nonsymmetric linear systems. For some nonsymmetric problems, letting restriction be the transpose of interpolation in a traditional AMG sense remains a robust solver [48]. However, at times this can produce coarse-grid matrices that are not amenable to recursive application of the method; for example, the diagonals could change sign or even be zero. Such situations were the original motivation for this work. There have also been efforts to develop AMG methods and underlying theory targeting general nonsymmetric matrices [13, 33, 39, 50, 56]. Nevertheless, *when* each of these approaches is appropriate requires experience and experimentation, which can limit their use and reliability outside of the AMG community.

In this work, we develop a novel, reduction-based AMG method to solve systems with triangular structure, with a particular focus on upwind discretizations of hyperbolic PDEs. Although geometric multigrid has been applied to upwind discretizations using the well-known line-smoother approach [43] and other hyperbolic-type problems [4, 57], an effective AMG algorithm does not need geometric information, a structured grid, or a fixed/known direction over which to relax. For these reasons, an effective AMG solver is much more robust and non-intrusive, in the sense that it can easily be used by software packages, regardless of their meshing, discretization, etc.. The model problem considered here is the steady-state transport equation, which comes up in large-scale simulation of neutron and radiation transport [2, 6, 16, 26, 44]. Although this work is focused on algorithm theory and development, the long-term goal is a highly-scalable parallel solver for upwind discretizations and other systems with triangular structure.

Results of our new algorithm significantly out-perform current state-of-the-art iterative solvers in the context of steady-state transport [52]. Even for high-order, discontinuous discretizations on unstructured grids, our method is able to achieve an order-of-magnitude residual reduction in 1–2 iterations. The method likely extends to other systems with triangular structure such as directed acyclic graphs, other hyperbolic PDEs, or certain full space-time discretizations of hyperbolic PDEs as well, but this is a topic for future study.

Background on AMG, along with theoretical motivation for the new method and its relation to existing algorithms, is discussed in Section 2. The underlying premise is that certain "ideal" operators can lead to a reduction-based AMG method; in particular, we develop an AMG algorithm based on approximating ideal restriction (AMGir). Section 3 introduces a natural way to approximate ideal operators for matrices with triangular structure, as well as theoretical results on the corresponding multigrid error propagation. The algorithm and possible variations are presented in Section 4, and numerical results and scaling studies in Section 5. Finally, conclusions and future directions of research will be discussed in Section 6.

2. Background and motivation.

2.1. Algebraic multigrid. Algebraic multigrid applied to the linear system $A\mathbf{x} = \mathbf{b}$, $A \in \mathbb{R}^{n\times n}$, consists of two processes: relaxation and coarse-grid correction, typically designed to be complementary in the sense that they reduce error associated with different parts of the spectrum. Relaxation takes the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k + M^{-1}(\mathbf{b} - A\mathbf{x}_k), \tag{1}$$

where M^{-1} is some approximation to A^{-1} such that the action of M^{-1} can be easily computed. For example, M may be the diagonal of A (Jacobi) or the upper or lower triangular block of A (Gauss-Seidel). Error associated with large eigenvalues of A, generally corresponding to geometrically high-frequency error that cannot be well-represented on a coarse grid, is targeted through relaxation. Coarse-grid correction is then built to reduce error not reduced by relaxation, typically algebraically smooth error, that is, modes associated with small eigenvalues. Coarse-grid correction takes the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k + P(RAP)^{-1}R(\mathbf{b} - A\mathbf{x}_k), \tag{2}$$

where $P \in \mathbb{R}^{n\times n_c}$ and $R \in \mathbb{R}^{n_c\times n}$ are interpolation and restriction operators, respectively, between $\mathbb{R}^n$ and the next coarser grid in the AMG hierarchy, $\mathbb{R}^{n_c}$.

A two-level $V(1,1)$-cycle is then given by combining the coarse-grid correction in (2) with pre- and post-relaxation steps as in (1), resulting in a two-grid error propagation operator of the form

$$E_{TG} = (I - M^{-1}A)(I - P(RAP)^{-1}RA)(I - M^{-T}A). \tag{3}$$

The goal in building an AMG solver is to bound $\|E_{TG}\| \ll 1$ in some norm, such that AMG iterations are guaranteed to converge quickly. For SPD problems, $\|\mathbf{x}\|_A^2 = \langle A\mathbf{x}, \mathbf{x}\rangle = \mathbf{x}^T A\mathbf{x}$ ($\langle\cdot,\cdot\rangle$ denoting the Euclidean inner product) defines a norm, referred to as the *A-norm* or *energy norm*. Typically, restriction is then defined as $R = P^T$, wherein the coarse-grid correction $P(P^TAP)^{-1}P^TA$ is exactly an A-orthogonal projection onto $\mathcal{R}(P)$, the range of P. For nonsymmetric systems considered here, $\langle A\mathbf{x}, \mathbf{x}\rangle$ no longer defines a valid norm. Various efforts have been made to develop a framework for convergence of nonsymmetric problems [13, 39], but choosing R such that coarse-grid correction is an orthogonal projection (in some inner-product) is often not feasible in practice. Thus, a restriction operator R must instead be chosen such that the oblique projection, $I - P(RAP)^{-1}RA$, has desirable properties for AMG convergence. Here, the concept of "ideal restriction" is theoretically motivated in Sections 2.2, followed by a practical approximation in Section 3 and an algorithm outline in Section 4.

The computational cost or complexity of an AMG algorithm is typically measured in *work units* (WU), where one WU is the cost to perform a sparse matrix-vector multiply on the finest level in the AMG hierarchy. This motivates two additional complexity measures in AMG, *operator complexity* (OC) and *cycle complexity* (CC). Operator complexity gives the cost in WUs to perform a sparse matrix-vector multiply on each level in the hierarchy, while CC gives the cost in WUs to perform one AMG iteration, including pre- and post-relaxation, computing and restricting the residual, and coarse-grid correction. Let A_ℓ, P_ℓ, and R_ℓ be operators on the ℓth level of the AMG hierarchy, and $|A_\ell|$ denote the number of non-zeros in matrix A_ℓ. Then

$$\text{OC} = \sum_\ell \frac{|A_\ell|}{|A_0|},$$

$$\mathrm{CC} = \sum_{\ell} \frac{(\nu_{\mathrm{pre}} + \nu_{\mathrm{post}} + 1)|A_\ell| + |P_\ell| + |R_\ell|}{|A_0|},$$

where ν_{pre} and ν_{post} are the number of pre- and post-relaxations, respectively, and $\ell = 0$ is the finest level in the hierarchy. Such measures lead to two additional objective measures of AMG performance, the *effective convergence factor* (ECF), ρ_{eff}, which measures the residual reduction factor for the equivalent of one WU, and *work-per-digit-of-accuracy*, χ_{wpd}, which measures the WUs necessary to achieve an order-of-magnitude reduction in the residual. For convergence factor ρ and cycle complexity CC,

$$\rho_{eff} := \rho^{\frac{1}{CC}}, \qquad \chi_{wpd} := -\frac{CC}{\log_{10}(\rho)} = -\frac{1}{\log_{10}(\rho_{eff})}.$$

We use these objective measures to motivate parts of the AMGir algorithm, and to consider the performance and scaling of AMGir in Section 5. Although these are good measures of serial performance, they do not reflect parallel efficiency of the algorithm.

2.2. Ideal restriction. In a classical AMG sense, suppose that a CF-splitting has been performed on the current grid, that is, all degrees-of-freedom (DOFs) have been designated as either a coarse-point (C-point) or a fine-point (F-point). C-points each correspond to a DOF on the next coarser grid in the AMG hierarchy, and are interpolated and restricted between grids by injection (i.e., by value) [10]. Let n_f denote the number of F-points and n_c the number of C-points, and, for example, let A_{ff} denote the block of F-to-F connections in A. Then, operators A, P, and R can be written in the following block forms:

$$A = \begin{pmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{pmatrix}, \quad P = \begin{pmatrix} W \\ I \end{pmatrix}, \quad R = \begin{pmatrix} Z & I \end{pmatrix}, \tag{4}$$

where $W \in \mathbb{R}^{n_f \times n_c}$ interpolates to F-points on the fine grid via linear combinations of coarse-grid DOFs, and $Z \in \mathbb{R}^{n_c \times n_f}$ restricts F-point residuals. Note that (4) implicitly assumes the same CF-splitting for R and P, although the sparsity patterns for nonzero elements of Z^T and W may be different. For notation, define the coarse-grid operator and oblique projection onto $\mathcal{R}(P)$, respectively, as

$$\begin{aligned} \mathcal{K} :&= RAP \\ &= ZA_{ff}W + A_{cf}W + ZA_{fc} + A_{cc}, \\ \pi &:= P(RAP)^{-1}RA. \end{aligned} \tag{5}$$

Here, $\mathcal{K}$ is used to denote the coarse-grid operator instead of the traditional notation, A_c, to avoid confusion with subscripts denoting C-points.

LEMMA 1 (Schur-complement coarse grid). *Let A take the block form as in (4) and A_{ff} be nonsingular. Define*

$$R_{ideal} := \begin{pmatrix} -A_{cf}A_{ff}^{-1} & I \end{pmatrix}, \quad P_{ideal} := \begin{pmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{pmatrix}.$$

Then, for all R and P of the form in (4),

$$RAP_{ideal} = R_{ideal}AP = R_{ideal}AP_{ideal} = S_A,$$

where $S_A := A_{cc} - A_{cf}A_{ff}^{-1}A_{fc}$ is the Schur-complement of A in A_{cc}. That is, the coarse-grid operator, $\mathcal{K}$, is given by the Schur complement of A.

Proof. Suppose $R = R_{ideal}$. Then, $Z = -A_{cf}A_{ff}^{-1}$ and

$$\begin{aligned} R_{ideal}AP &= ZA_{ff}W + A_{cf}W + ZA_{fc} + A_{cc} \\ &= -A_{cf}W + A_{cf}W - A_{cf}A_{ff}^{-1}A_{fc} + A_{cc} \\ &= S_A. \end{aligned}$$

Identical steps show the equivalent result for P_{ideal}. ∎

The operator R_{ideal} defined in Lemma 1 is referred to here as "ideal restriction," a natural extension of the well-known "ideal interpolation" operator, P_{ideal}, also defined in Lemma 1. Note that if A is symmetric, $R_{ideal} = P_{ideal}^T$. Ideal interpolation is well-motivated under classical AMG theory (for example, see [22]). Here, we show that ideal restriction and ideal interpolation each have significance in the nonsymmetric setting as well. In Theorem 2, ideal restriction is shown to be ideal in a certain sense, albeit not the same sense for which ideal interpolation acquired its name. Thus, consider error-propagation for a coarse-grid correction applied to the current error vector, $(\mathbf{e}_f, \mathbf{e}_c)^T$:

$$\begin{aligned} (I-\pi)\mathbf{e} &= \left[\begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix} - \begin{pmatrix} W \\ I \end{pmatrix} \mathcal{K}^{-1} \begin{pmatrix} Z & I \end{pmatrix} \begin{pmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{pmatrix} \right] \begin{pmatrix} \mathbf{e}_f \\ \mathbf{e}_c \end{pmatrix} \\ &= \begin{pmatrix} I - W\mathcal{K}^{-1}(ZA_{ff} + A_{cf}) & -W\mathcal{K}^{-1}(ZA_{fc} + A_{cc}) \\ -\mathcal{K}^{-1}(ZA_{ff} + A_{cf}) & I - \mathcal{K}^{-1}(ZA_{fc} + A_{cc}) \end{pmatrix} \begin{pmatrix} \mathbf{e}_f \\ \mathbf{e}_c \end{pmatrix} \\ &= \begin{pmatrix} \Big[I - W\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\Big]\mathbf{e}_f - W\mathcal{K}^{-1}(ZA_{fc} + A_{cc})\mathbf{e}_c \\ \Big[I - \mathcal{K}^{-1}(ZA_{fc} + A_{cc})\Big]\mathbf{e}_c - \mathcal{K}^{-1}(ZA_{ff} + A_{cf})\mathbf{e}_f \end{pmatrix}. \end{aligned} \tag{6}$$

Now, let $\mathbf{e}_f = W\mathbf{e}_c + \delta_f$, that is, let $\mathbf{e}_f$ be interpolated from $\mathbf{e}_c$ with some error term δ_f, and recall $\mathcal{K} = ZA_{ff}W + A_{cf}W + ZA_{fc} + A_{cc}$. Then, plugging $W\mathbf{e}_c + \delta_f$ into the fine-grid block, and $ZA_{fc} + A_{cc} = \mathcal{K} - ZA_{ff}W - A_{cf}W$ into the coarse-grid block gives:

$$\begin{aligned} (I-\pi)\mathbf{e} &= \begin{pmatrix} \Big[I - W\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\Big](W\mathbf{e}_c + \delta_f) - W\mathcal{K}^{-1}(ZA_{fc} + A_{cc})\mathbf{e}_c \\ \mathcal{K}^{-1}(ZA_{ff} + A_{cf})(W\mathbf{e}_c - \mathbf{e}_f) \end{pmatrix} \\ &= \begin{pmatrix} \Big[I - W\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\Big]\delta_f + W\mathbf{e}_c - W\mathcal{K}^{-1}\mathcal{K}\mathbf{e}_c \\ -\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\delta_f \end{pmatrix} \\ &= \begin{pmatrix} \Big[I - W\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\Big]\delta_f \\ -\mathcal{K}^{-1}(ZA_{ff} + A_{cf})\delta_f \end{pmatrix}. \end{aligned} \tag{7}$$

The following theorem follows from (7).

THEOREM 2 (Ideal restriction). *For a given CF-splitting, assume that A_{ff} is full rank and let A take the block form as given in (4). Assume that C-points are interpolated and restricted by injection in a classical-AMG sense (4). Then, an exact coarse-grid correction at C-points is attained for all $\mathbf{e}$ if and only if*

$$R = R_{ideal} := \begin{pmatrix} -A_{cf}A_{ff}^{-1} & I \end{pmatrix}. \tag{8}$$

Furthermore, the error in the coarse-grid correction is given by

$$(I-\pi)\mathbf{e} = \begin{pmatrix} \mathbf{e}_f - W\mathbf{e}_c \\ \mathbf{0} \end{pmatrix}, \tag{9}$$

where $P = \begin{pmatrix} W \\ I \end{pmatrix}$. *Finally, a coarse-grid correction using* R_{ideal} *followed by an exact solve on F-points results in an exact two-grid method, independent of* W.

Proof. From (7), coarse-grid correction as applied to some vector $\mathbf{e} = (\mathbf{e}_f, \mathbf{e}_c)^T$ and restricted to C-points is given by $\mathcal{K}^{-1}(ZA_{ff}+A_{cf})(W\mathbf{e}_c-\mathbf{e}_f)$. Noting that there does not exist a W such that $W\mathbf{e}_c = \mathbf{e}_f$ for all vectors $\mathbf{e} = (\mathbf{e}_f, \mathbf{e}_c)^T$, it follows that $\mathcal{K}^{-1}(ZA_{ff} + A_{cf})(W\mathbf{e}_c - \mathbf{e}_f) = 0$ if and only if $ZA_{ff} + A_{cf} = 0$. Given that A_{ff} is nonsingular, $ZA_{ff} + A_{cf} = 0$ if and only if $Z = -A_{cf}A_{ff}^{-1}$. The error shown in (9) follows directly from plugging R_{ideal} into (7). An exact solve on F-points would eliminate this error, providing an exact two-grid method. ∎

COROLLARY 3 (Ideal interpolation). *An exact solve on F-points followed by a coarse-grid correction using so-called "ideal interpolation,"* $P := P_{ideal}$, *as defined in Lemma 1 gives an exact two-level method, independent of* Z.

Proof. An initial exact solve on F-point rows of the equation $A\mathbf{x} = \mathbf{b}$ means that the residual at F-point rows is zero. For given error vector $\mathbf{e}$, this is equivalent to updating $\mathbf{e}$ such that $A_{ff}\mathbf{e}_f + A_{fc}\mathbf{e}_c = 0$, or $\mathbf{e}_f = -A_{ff}^{-1}A_{fc}\mathbf{e}_c$. Plugging this into the coarse-grid correction expansion in (6), along with $W = -A_{ff}^{-1}A_{fc}$ and $\mathcal{K} = S_A$ (see Lemma 1) gives

$$\begin{aligned}(I-\pi)\mathbf{e} &= \begin{pmatrix} -A_{ff}^{-1}A_{fc}\mathbf{e}_c \\ \mathbf{e}_c \end{pmatrix} - \begin{pmatrix} -A_{ff}^{-1}A_{fc} \\ I \end{pmatrix} \mathcal{S}_{\mathcal{A}}{}^{-1} \begin{pmatrix} Z & I \end{pmatrix} \begin{pmatrix} \mathbf{0} \\ S_A\mathbf{e}_c \end{pmatrix} \\ &= \mathbf{0}. \end{aligned}$$ ∎

It follows from Theorem 2 and Corollary 3 that R_{ideal} and P_{ideal} each lead to an exact two-level method, independent of the accompanying interpolation and restriction operators, respectively, when coupled with an exact solve on F-points. Such a two-level method is in fact a reduction algorithm, as solving $A\mathbf{x} = \mathbf{b}$ is reduced to solving two smaller systems. Note that the ordering of solving the coarse- and fine-grid problems is fundamental to achieving reduction. That is, to achieve reduction, the F-point solve must *follow* coarse-grid correction with R_{ideal}, while the F-point solve must *precede* coarse-grid correction with P_{ideal}. Background on reduction algorithms and ideal restriction in the multigrid context is given in Sections 2.3 and 3.

A natural corollary of Theorem 2 is that if we define "trivial interpolation" as $P_0 = \begin{pmatrix} 0 \\ I \end{pmatrix}$, then error propagation of coarse-grid correction with R_{ideal} and P_0 is an ℓ^2-orthogonal projection. Similarly, if we define "trivial restriction" as $R_0 = \begin{pmatrix} 0 & I \end{pmatrix}$, then residual propagation of coarse-grid correction with R_0 and P_{ideal} is an ℓ^2-orthogonal projection. This is consistent with the proofs of Theorem 2 and Corollary 3, that is, P_{ideal} is based on eliminating residuals and R_{ideal} on eliminating error. Despite orthogonality, in practice, choosing $W = 0$ for interpolation of F-points does not always provide the best convergence factors. This is discussed in greater detail in Section 4.6.

2.3. Relation to existing AMG methods. The goal of a reduction-based algorithm is simply to take a problem and split its solving into smaller sub-problems. One way to enable reduction is through block row-elimination, where A can be written as

$$\begin{pmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{pmatrix} = \begin{pmatrix} I & 0 \\ A_{cf}A_{ff}^{-1} & I \end{pmatrix} \begin{pmatrix} A_{ff} & 0 \\ 0 & S \end{pmatrix} \begin{pmatrix} I & A_{ff}^{-1}A_{fc} \\ 0 & I \end{pmatrix}. \tag{10}$$

Using (10), solving $A\mathbf{x} = \mathbf{b}$ is reduced to solving two systems of size $n_F \times n_F$ and $n_C \times n_C$, the latter of which is the Schur complement, $S := A_{cc} - A_{cf}A_{ff}^{-1}A_{fc}$. Unfortunately, (10) is generally impractical as shown because of the need to compute the inverse of A_{ff} and solve the typically dense Schur-complement system. However, such block decompositions have been used to motivate many algorithms, and a new approximation to inverting (10) can be found in Section 3.3.

The goal of this paper is to develop a reduction-based AMG algorithm for highly nonsymmetric linear systems. However, Schur-complement coarse-grid operators and reduction-based multigrid are not new to the multigrid literature. Reduction-based multigrid has been considered in many contexts, starting in the geometric multigrid setting [24, 47]. A high-performance variation of geometric multigrid based on a reduction point-of-view is the SMG solver in the Hypre library, which uses a sparse approximation to P_{ideal} that reproduces the action of P_{ideal} on the constant vector [23]. The SMG solver, however, assumes knowledge of an underlying structured grid, and forms symmetric, Galerkin coarse grids. More recently, an adaptive, two-level reduction-based AMG method was proposed in [11, 30], where an approximation to P_{ideal} is developed through an adaptive process. This work provided interesting theoretical results; however, they are based in the A-norm, and thus again confined to the SPD setting. As seen in Section 2.2, a reduction-based algorithm is invariably linked to a Schur-complement coarse-grid operator. Other literature exists that is not explicitly marketed as reduction-based, but targets a Schur-complement coarse grid, for example, [15, 31, 34, 37, 38, 49]. Finally, the multigrid reduction-in-time parallel-in-time algorithm [18, 19, 20] is also based on (10), and in [20, 25] it was shown that the well-known *parareal* parallel-in-time method [25] is equivalent to a two-level multigrid reduction-in-time method.

Algebraic multigrid methods have also been developed that target nonsymmetric systems, but the theory is limited and the development of effective solvers is ongoing. The reduction-based decomposition given in (10) was considered in [38] for nonsymmetric M-matrices resulting from convection-diffusion discretizations. An AMG algorithm for nonsymmetric systems was proposed in [56] based on (10), where it is shown that, coupled with an exact solve on F-points, convergence in the spectral sense is bounded based on how well the coarse grid approximates the Schur complement.

Ideal restriction and interpolation are approximated in [56] by performing a constrained minimization over a fixed sparsity pattern for R and P, where known near-null space modes are interpolated exactly, and the remaining DOFs used to approximate the ideal operators. Such an approach is similar to the constrained minimization approach for nonsymmetric systems used in [33, 42], which approximates the ideal operators of A^*A and AA^* for P and R, respectively. Although reduction is not achieved for ideal operators based on A^*A and AA^*, such operators are well-motivated under traditional AMG convergence theory [33]. In fact, in the symmetric setting, approximating ideal operators for A vs. A^2 through constrained energy-minimization produce nearly identical performance [33]. As an extension to the nonsymmetric setting, it is hypothesized that the AMG solvers in [33, 56] would achieve similar performance, although such a comparison has not yet been done. Regardless, the constrained approximation of ideal operators such as in [33, 56] is a broadly applicable solver that relies heavily on constraints for good convergence [33, 53].

This leads to the novelty of this work. Here, we show that matrices with triangular structure are amenable to a reduction-type algorithm (Section 3), and we develop a natural way to approximate ideal operators for triangular matrices. The resulting solver does not rely on constraints for strong convergence, and is based only

on an appropriate CF-splitting and the approximation of R_{ideal}, which are discussed in Section 3.

3. Triangular structure and approximate ideal restriction. In Section 2.2, theory was developed to motivate ideal restriction and ideal interpolation for nonsymmetric linear systems. However, ideal operators are typically not formed in practice due to the complexity of forming A_{ff}^{-1}. This section shows that for block-triangular matrices, there is a natural way to approximate A_{ff}^{-1}. Furthermore, approximating ideal operators in coarse-grid correction, coupled with Jacobi F-relaxation, gives a nilpotent multigrid error-propagation operator, that is, convergence to the exact solution is ultimately guaranteed (Section 3.1). Coupled with an appropriate CF-splitting (Section 3.2), this leads to a fast and practical solver for matrices with triangular structure.

Although direct solves of triangular matrices do not scale well in parallel, direct solve considerations suggest that a triangular matrix is amenable to a reduction-based algorithm, because each step in a forward or backward solve is effectively reduction by eliminating one DOF. Thus, the goal here is to develop a reduction-based solver for triangular systems, independent of matrix ordering, based on the concept of "approximate ideal restriction" (AIR). Moving forward, assume that A is lower triangular with unit diagonal in some ordering. For theoretical purposes, we assume that A is actually ordered to be lower triangular, but it is important to note that results presented here are independent of this ordering, and it is only used to simplify proofs. An in-depth discussion on handling of block-triangular matrices is given in Section 4.1.

Let $A_{ff} = I - L_{ff}$, where L_{ff} is the strictly lower triangular part of A_{ff}, and is, thus, also nilpotent. Hence, A_{ff}^{-1} can be written as a finite Neumann expansion:

$$A_{ff}^{-1} = (I - L_{ff})^{-1} = \sum_{i=0}^{n} L_{ff}^{i}. \tag{11}$$

An order-k approximation to A_{ff}^{-1} is then given by truncating (11): $\Delta^{(k)} := \sum_{i=0}^{k} L_{ff}^{i}$, for some $0 \leq k \leq n$. The error in $\Delta^{(k)}$ can be measured as $I - \Delta^{(k)} A_{ff} = L_{ff}^{k+1}$. Noting that $R_{\text{ideal}} A = \begin{pmatrix} 0 & S \end{pmatrix}$, let $R = \begin{pmatrix} -A_{cf}\Delta^{(k)} & I \end{pmatrix}$ and consider its action on A:

$$\begin{aligned} RA &= \begin{pmatrix} A_{cf}(I - \Delta^{(k)} A_{ff}) & A_{cc} - A_{cf}\Delta^{(k)} A_{fc} \end{pmatrix} \\ &= \begin{pmatrix} A_{cf} L_{ff}^{k+1} & A_{cc} - A_{cf}\Delta^{(k)} A_{fc} \end{pmatrix}. \end{aligned}$$

Since L_{ff} is nilpotent, approximating ideal restriction via a truncated Neumann expansion of A_{ff}^{-1} can be seen as equivalent to approximating the action of R_{ideal} on A, namely, trying to set RA equal to zero within F-point columns. For certain structured matrices, the kth-order Neumann expansion is exactly eliminating the contribution of F-points within distance k for a given C-point (row of R). In fact, this was the original motivation for AIR – eliminating the contribution of error at F-points to the coarse-grid right-hand side.

Numerical results presented in Section 5 are based on two properties of the AMGir error-propagation operator. In a reduction context, the error-propagation operator in the multilevel setting is nilpotent, ensuring an (eventually) exact method (Section 3.1). In addition to the nilpotency, an appropriate CF-splitting such that off-diagonal elements of A_{ff} are small additionally ensures effective error reduction at indices that

are not "reduced" in a given iteration (Section 3.2). The latter is likely the primary reason for strong convergence, although both are important.

3.1. Nilpotency of error propagation. Lemma 4 and Corollary 5 below are presented as important pieces in motivating AMGir. Corollary 5 is fundamental to a multilevel implementation of AMGir, showing that as long as A_{ff}^{-1} in ideal restriction and interpolation is approximated with a lower-triangular matrix, the coarse-grid operator retains the same triangular structure as the fine grid. Thus, the same algorithms for forming P and R are appropriate in a recursive and multilevel fashion. For given ideal restriction and interpolation approximations, coupled with Jacobi relaxation on F-points, Lemma 6 shows that the two-grid error propagation operator is nilpotent. Theorem 7 is the primary convergence result of this section, showing that a full multilevel implementation has a nilpotent error-propagation operator. A trivial corollary is that convergence to the exact solution is guaranteed within n iterations. Although bounds of $O(n)$ iterations of $O(n)$ complexity is far removed from the typical linear or near-linear complexity of AMG, such a bound is at least of theoretical interest.

LEMMA 4. *Let A be lower triangular and suppose the DOFs of A have been partitioned into C-points and F-points. Then, if $\Delta \in \mathbb{R}^{n_f \times n_f}$ is lower triangular, $A_{cf}\Delta A_{fc}$ is nilpotent and strictly lower triangular. The same holds for $A_{fc}\hat{\Delta}A_{cf}$, where $\hat{\Delta} \in \mathbb{R}^{n_c \times n_c}$.*

Proof. Let C-points be given by $\mathcal{C} = \{\mathcal{C}_1, ..., \mathcal{C}_{n_c}\}$, where $\mathcal{C}_i$ denotes the global index of the ith C-point, and likewise for F-points. Furthermore, assume $\mathcal{C}$ and $\mathcal{F}$ are in increasing order, that is, $\mathcal{C}_1 < \mathcal{C}_2 < ... < \mathcal{C}_{n_c}$, and that global indices of A are ordered such that A is lower triangular, that is, $A_{ij} = 0$ if $i \leq j$. In terms of paths in the graph of A, this means that from a given node i, there only exist paths to nodes $j \leq i$.

For $A_{cf}\Delta A_{fc}$ to be strictly lower triangular, we must have $(A_{cf}\Delta A_{fc})_{jk} = 0$ for all $0 \leq j \leq k \leq n_c$. This can be shown by considering paths between the jth and kth C-points in the graph of $A_{cf}\Delta A_{fc}$.

1. Starting at $\mathcal{C}_j$, A_{fc} can only contain paths from $\mathcal{C}_j$ to F-points $\mathcal{F}_i < \mathcal{C}_j$.
2. Then, Δ can only contain paths from $\mathcal{F}_i$ to F-points $\mathcal{F}_\ell \leq \mathcal{F}_i$.
3. Finally, A_{cf} maps from F-points $\mathcal{F}_\ell$ to C-points $\mathcal{C}_k < \mathcal{F}_\ell$.

Thus, the only possible paths in the graph of $A_{cf}\Delta A_{fc}$ are from points $\mathcal{C}_j$ to $\mathcal{C}_k < \mathcal{C}_j$. It follows that if $\mathcal{C}_k \geq \mathcal{C}_j$, or, equivalently, $k \geq j$, then $(A_{cf}\Delta A_{fc})_{jk} = 0$. ∎

COROLLARY 5. *Let A be lower triangular and Δ_R and Δ_P be lower triangular approximations to A_{ff}^{-1}, and define $P := \begin{pmatrix} -\Delta_P A_{fc} \\ I \end{pmatrix}$ and $R := \begin{pmatrix} -A_{cf}\Delta_R & I \end{pmatrix}$. Then, the coarse-grid operator $\mathcal{K} := RAP$ (5) is lower triangular with diagonal given by that of A_{cc}.*

Proof. First note that

$$\mathcal{K} := RAP = A_{cc} - A_{cf}(\Delta_R + \Delta_P - \Delta_R A_{ff} \Delta_P)A_{fc}.$$

We know that A_{cc} is lower triangular. It follows from Lemma 4 that $A_{cf}(\Delta_R + \Delta_P - \Delta_R A_{ff}\Delta_P)A_{fc}$ is strictly lower triangular. Thus, $\mathcal{K}$ is lower triangular with the same diagonal as A_{cc}. ∎

LEMMA 6 (Nilpotent two-grid AMG). *Let A be lower triangular in some ordering and Δ_R and Δ_P be lower triangular approximations to A_{ff}^{-1}, and define $P :=$*

$\begin{pmatrix} -\Delta_P A_{fc} \\ I \end{pmatrix}$ *and* $R := \begin{pmatrix} -A_{cf}\Delta_R & I \end{pmatrix}$. *Then, the two-grid AMG preconditioner with coarse-grid correction and Jacobi F-relaxation is strictly lower triangular, and thus nilpotent.*

Proof. Define C-points as the set $\mathcal{C} = \{\mathcal{C}_1, ..., \mathcal{C}_{n_c}\}$, where $\mathcal{C}_i$ is the global index of the ith C-point, and likewise for F-points. Furthermore, let $\mathcal{C}$ and $\mathcal{F}$ be in increasing order, that is, $\mathcal{C}_1 < \mathcal{C}_2 < ... < \mathcal{C}_{n_c}$, and let global indices of A be ordered such that A is lower triangular, that is, $A_{ij} = 0$ if $i < j$. Then, for a given node i in A, there only exists paths to nodes $j \leq i$. It follows that off-diagonal block A_{cf} only maps from F_i to $C_j < F_i$, and similarly for A_{fc}.

The two-grid AMG error-propagation operator with relaxation scheme M^{-1} looks like $E_{TG} = (I - M^{-1}A)(I - P(RAP)^{-1}RA)$. First consider the coarse-grid correction term:

$$\begin{aligned} I - P(RAP)^{-1}RA &= \begin{pmatrix} I & 0 \\ 0 & I \end{pmatrix} - \begin{pmatrix} -\Delta_P A_{fc} \\ I \end{pmatrix} \mathcal{K}^{-1} \begin{pmatrix} -A_{cf}\Delta_R & I \end{pmatrix} \begin{pmatrix} A_{ff} & A_{fc} \\ A_{cf} & A_{cc} \end{pmatrix} \\ &= \begin{pmatrix} I + \Delta_P A_{fc}\mathcal{K}^{-1}A_{cf}(I - \Delta_R A_{ff}) & \Delta_P A_{fc}\mathcal{K}^{-1}(A_{cc} - A_{cf}\Delta_R A_{fc}) \\ -\mathcal{K}^{-1}A_{cf}(I - \Delta_R A_{ff}) & I - \mathcal{K}^{-1}(A_{cc} - A_{cf}\Delta_R A_{fc}) \end{pmatrix}. \end{aligned} \tag{12}$$

Note that $I - \Delta_R A_{ff}$ is strictly lower triangular, and $\mathcal{K}$ and $A_{cc} - A_{cf}\Delta_R A_{fc}$ are lower triangular. Then, the action of the lower-left block in (12) can be seen as three steps, mapping $F_i \to F_j \to C_k \to C_\ell$, where $C_\ell \leq C_k < F_j < F_i$. A similar result holds for the upper-right block, showing that, in the global ordering, the off-diagonal blocks of (12) are strictly lower triangular. For the lower-right block, it follows from Corollary 5 and Lemma 4 that $\mathcal{K}^{-1}(A_{cc} - A_{cf}\Delta_R A_{fc})$ is lower triangular with unit diagonal. Subtracting from the identity gives a strictly lower triangular block. Finally, given that the product of a strictly lower triangular matrix and lower triangular matrix is strictly lower triangular, it follows from Lemma 4 that $\Delta_P A_{fc}\mathcal{K}^{-1}A_{cf}(I - \Delta_R A_{ff})$ is strictly lower triangular. Then, the upper left block is lower triangular with unit diagonal.

Thus, the error-propagation operator for coarse-grid correction is lower triangular with unit diagonal in the FF-block, and strictly lower triangular in other blocks. Furthermore, the error-propagation matrix for Jacobi relaxation on F-points is strictly lower triangular in the FF-block and lower triangular elsewhere. It follows that the AIR two-grid error-propagation operator given by the product of F-Jacobi relaxation and coarse-grid correction is a strictly lower triangular and, thus, nilpotent, operator. ∎

THEOREM 7 (Nilpotent AMG). *Let A be lower triangular in some ordering and Δ_R and Δ_P be lower triangular approximations to A_{ff}^{-1}, and define $P := \begin{pmatrix} -\Delta_P A_{fc} \\ I \end{pmatrix}$ and $R := \begin{pmatrix} -A_{cf}\Delta_R & I \end{pmatrix}$. Then, the multilevel AMG preconditioner with coarse-grid correction and Jacobi F-relaxation is strictly lower triangular, and thus nilpotent.*

Proof. In proving a nilpotent error-propagation for a multilevel algorithm, we proceed in a recursive fashion. Consider the error-propagation of coarse-grid correction with an inexact coarse-grid solve. Error propagation of coarse-grid correction takes the form $\mathbf{e}_{new} = \mathbf{e}_{old} - P\mathbf{x}_c$, for some coarse-grid correction $\mathbf{x}_c$. In the case of an exact coarse-grid solve, $P\overline{\mathbf{x}}_c = P\mathcal{K}^{-1}RA\mathbf{e}_{old}$ and $\mathbf{e}_{new} = (I - P\mathcal{K}^{-1}RA)\mathbf{e}_{old}$. For an inexact coarse-grid solve, the correction can be written as some perturbation of the exact solve, $\mathbf{x}_c = \overline{\mathbf{x}}_c + (\mathbf{x}_c - \overline{\mathbf{x}}_c)$, where the error in the correction is given by $\mathbf{e}_c = \mathbf{x}_c - \overline{\mathbf{x}}_c$. The error propagation of the inexact coarse-grid solve, that is, the

coarse-grid V-cycle, then operates on $\mathbf{e}_c$, that is, $E_c(\mathbf{e}_c) = E_c(\mathbf{x}_0 - \overline{\mathbf{x}})$, for some initial guess $\mathbf{x}_0$. Assuming a zero initial guess on the coarse-grid problem as is standard in AMG, then

$$\begin{aligned}\mathbf{e}_{new} &= \mathbf{e}_{old} - P(\overline{\mathbf{x}}_c - E_c\overline{\mathbf{x}}_c)\\ &= \mathbf{e}_{old} - P(I - E_c)\overline{\mathbf{x}}_c\\ &= \mathbf{e}_{old} - P(I - E_c)\mathcal{K}^{-1}RA\mathbf{e}_{old}\\ &= \Big[(I - P\mathcal{K}^{-1}RA) + PE_c\mathcal{K}^{-1}RA\Big]\mathbf{e}_{old}.\end{aligned}$$

Lemma 6 showed that $(I - P\mathcal{K}^{-1}RA)$ is strictly lower triangular when coupled with Jacobi F-relaxation. Now, we must show that the error term $PE_c\mathcal{K}^{-1}RA$ is also strictly lower triangular. Given that A is lower triangular, it is sufficient to show that $PE_c\mathcal{K}^{-1}R$ is strictly lower triangular. Expanding, we get

$$PE_c\mathcal{K}^{-1}R = \begin{pmatrix}\Delta_P A_{fc}E_c\mathcal{K}^{-1}A_{cf}\Delta_R & -\Delta_P A_{fc}E_c\mathcal{K}^{-1}\\ -E_c\mathcal{K}^{-1}A_{cf}\Delta_R & E_c\mathcal{K}^{-1}\end{pmatrix}. \tag{13}$$

An analogous result to Lemma 6 confirms that if $E_c\mathcal{K}^{-1}$ is strictly lower triangular, then so is (13).

Let E_i be the error-propagation operator for a V-cycle starting on the ith level of the hierarchy. On the coarsest level of the hierarchy, say $\ell = 0$, $E_0 = \mathbf{0}$ because the solve is exact. Then, on level $\ell = 1$, $E_c\mathcal{K}^{-1} = E_0\mathcal{K}^{-1} = \mathbf{0}$ is strictly lower triangular. It follows that the error-propagation operator on level $\ell = 1$, E_1, is strictly lower-triangular. On level $\ell = 2$, $E_c\mathcal{K}^{-1} = E_1\mathcal{K}^{-1}$ is again strictly lower triangular, because E_1 is strictly lower triangular and $\mathcal{K}$ lower triangular. By induction, the error-propagation operator E_ℓ for a multilevel hierarch with ℓ levels is strictly lower triangular and, thus, nilpotent. ∎

For triangular systems, it is easy to see that simple Jacobi relaxation also produces a nilpotent error-propagation operator. The significance of Lemma 6 and Theorem 7 is in showing analytically that (at least in the spectral sense) coarse-grid correction does not cause divergent behavior, and AMGir is asymptotically robust for triangular systems.

3.2. CF-splitting and L_{ff}. Although nilpotency of error-propagation is a nice property to have, it does not necessarily imply fast convergence. To understand what else is contributing to good convergence, we further examine the error-propagation operator. First note that error propagation for Jacobi F-relaxation is given by $I - D_{ff}^{-1}A_{ff} = L_{ff}$. Letting $\Delta_P = \Delta_R := \Delta^{(k)}$ for some $k \geq 0$, then the following identities hold, which we will need later:

$$\begin{aligned}\mathcal{K} &= A_{cc} - A_{cf}\Delta^{(2k+1)}A_{fc},\\ \left(\Delta^{(2k+1)}\right)^{-1} &= \left(I + L_{ff}^{k+1}\right)^{-1}\left(\Delta^{(k)}\right)^{-1},\\ \left(\Delta^{(2k+1)}\right)^{-1}\Delta^{(k)} &= I + \sum_{j=1}^{n}\left(-L_{ff}^{k+1}\right)^j.\end{aligned} \tag{14}$$

For ease of notation, let $\Delta := \Delta^{(k)}$ and $\tilde{\Delta} := \Delta^{(2k+1)}$. Then, using (14) and a Woodbury inverse expansion for $\mathcal{K}^{-1}$, we get

$$\mathcal{K}^{-1}(A_{cc} - A_{cf}\Delta A_{fc}) = \left(A_{cc}^{-1} + A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}A_{fc}A_{cc}^{-1}\right)(A_{cc} - A_{cf}\Delta A_{fc})$$

$$
\begin{aligned}
&= \left(I + A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}A_{fc}\right)\left(I - A_{cc}^{-1}A_{cf}\Delta A_{fc}\right)\\
&= I + A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}(I - \tilde{\Delta}^{-1}\Delta)A_{fc}\\
&= I - A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}\sum_{j=1}^{n}\left(-L_{ff}^{k+1}\right)^{j}A_{fc}. \qquad (15)
\end{aligned}
$$

Plugging (15) and the identity $I - \Delta^{(k)}A_{ff} = L_{ff}^{k+1}$ into the coarse-grid correction (12) shows that error propagation for coarse-grid correction can be written as

$$
\begin{aligned}
\mathbf{e}_f^{(1)} &= \left(I + \Delta A_{fc}\mathcal{K}^{-1}A_{cf}L_{ff}^{k+1}\right)\mathbf{e}_f^{(0)} +\\
&\qquad \Delta A_{fc}\left(I - A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}\sum_{j=1}^{n}\left(-L_{ff}^{k+1}\right)^{j}A_{fc}\right)\mathbf{e}_c^{(0)},\\
\mathbf{e}_c^{(1)} &= -\mathcal{K}^{-1}A_{cf}L_{ff}^{k+1}\mathbf{e}_f^{(0)} + A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}\sum_{j=1}^{n}\left(-L_{ff}^{k+1}\right)^{j}A_{fc}\mathbf{e}_c^{(0)}.
\end{aligned}
$$

Define $\mathcal{A} := \Delta A_{fc}$, $\mathcal{B} := \mathcal{K}^{-1}A_{cf}$, and $\mathcal{C} := A_{cc}^{-1}A_{cf}\left(\tilde{\Delta}^{-1} - A_{fc}A_{cc}^{-1}A_{cf}\right)^{-1}$. Then, if we consider the highest-order term of L_{ff}, error propagation takes the form

$$
\begin{aligned}
\mathbf{e}_f^{(i+1)} &= \left(I + \mathcal{A}\mathcal{B}L_{ff}^{k+1}\right)\mathbf{e}_f^{(i)} + \mathcal{A}\left(I - \mathcal{C}L_{ff}^{k+1}A_{fc}\right)\mathbf{e}_c^{(i)},\\
\mathbf{e}_c^{(i+1)} &= -\mathcal{B}L_{ff}^{k+1}\mathbf{e}_f^{(i)} + \mathcal{C}L_{ff}^{k+1}A_{fc}\mathbf{e}_c^{(i)}. \qquad (16)
\end{aligned}
$$

The key thing to notice in (16) is that each error update on C-points is hit with L_{ff}^{k+1}. Note that ℓ iterations of F-Jacobi relaxation has error propagation

$$
\begin{pmatrix}\mathbf{e}_f^{(i+\ell)}\\ \mathbf{e}_c^{(i+\ell)}\end{pmatrix} = \begin{pmatrix}L_{ff}^{\ell} & -\Delta^{(\ell-1)}A_{fc}\\ 0 & I\end{pmatrix}\begin{pmatrix}\mathbf{e}_f^{(i)}\\ \mathbf{e}_c^{(i)}\end{pmatrix}. \qquad (17)
$$

Then, coupling coarse-grid correction (16) with $k+1$ iterations of pre- or post-F-relaxation scales the error at *all* points by L_{ff}^{k+1}.

Such a scaling of error by L_{ff}^{k+1} is a complementary contribution to strong convergence factors achieved by AMGir. Ensuring that L_{ff} is not only nilpotent, but also has very small elements, is key to achieving a fast reduction in error. Here, we focus on matrices that result from the discretization of differential operators, in which case off-diagonal entries are typically expected to be smaller in magnitude than the diagonal. In fact, the traditional AMG motivation for a CF-splitting, picking F-points such that A_{ff} is well-conditioned, is also appropriate for AMGir. A classical AMG coarsening approach targets A_{ff} to be strictly diagonally dominant, which is equivalent to minimizing the size of elements in L_{ff}. Figure 1 demonstrates that for a model problem introduced in Section 5, L_{ff} has very small row sums, although very few rows of A_{ff} are diagonal (in which case the row of L_{ff} is empty). Similar results hold on all levels of the hierarchy, as well as for other discretizations and finite element orders.

Small row sums of L_{ff} reduce error efficiently for most rows. For rows without small row sums in L_{ff}, we rely on the redistribution of error through operators $\mathcal{A}, \mathcal{B}$, and $\mathcal{C}$, as well as the nilpotency of error propagation to reduce error. Specifically, larger row sums in L_{ff} typically occur near the boundary of the domain. Near the boundary, a nilpotent error-propagation operator reaches the exact solution in $O(1)$ iterations.

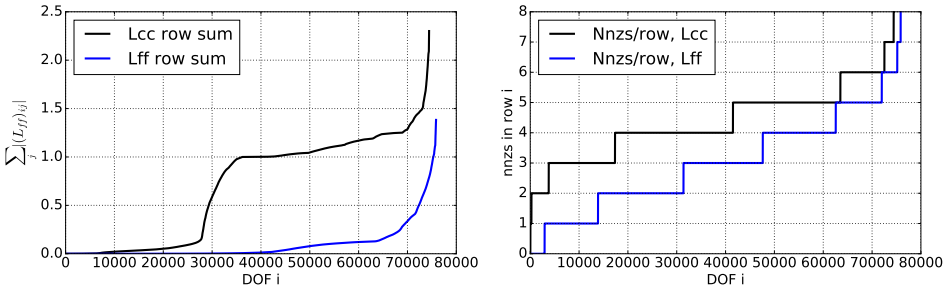

(a) Absolute row sum

(b) Nonzeros per row

Fig. 1: Data on CF-splitting and L_{ff} for the first level in a hierarchy built on a second-order finite element discretization of the *inset* model problem (see Section 5).

3.3. Schur-complement preconditioning. This work focuses on a traditional AMG framework. However, preconditioners are often developed by approximating the block decomposition of A in (10). An analogous result to Theorem 7 holds for the block preconditioner

$$\widehat{M} = \begin{pmatrix} I & 0 \\ A_{cf}\Delta_R & I \end{pmatrix} \begin{pmatrix} \Delta^{-1} & 0 \\ 0 & \widehat{\mathcal{K}} \end{pmatrix} \begin{pmatrix} I & \Delta_P A_{fc} \\ 0 & I \end{pmatrix}, \tag{18}$$

where Δ, Δ_R, and Δ_P are lower-triangular approximations to A_{ff}^{-1} and $\widehat{\mathcal{K}}$ some approximation to the Schur complement (not necessarily RAP), that is, the multilevel error-propagation operator, $I - \widehat{M}^{-1}A$, is nilpotent. Letting $\Delta = \Delta_P = \Delta_R$ be an order-k truncated Neumann expansion as in (11) and $\widehat{\mathcal{K}} := A_{cc} - A_{cf}\Delta A_{fc}$ gives the relation $A = \widehat{M} + N$, where $N = \begin{pmatrix} A_{ff} - \Delta^{-1} & 0 \\ 0 & 0 \end{pmatrix}$. Some linear algebra then shows that

$$I - \widehat{M}^{-1}A = \begin{pmatrix} I + \Delta A_{fc}\mathcal{K}^{-1}A_{cf} & 0 \\ -\mathcal{K}^{-1}A_{cf} & 0 \end{pmatrix} \begin{pmatrix} (D_{ff}^{-1}L_{ff})^{k+1} & 0 \\ 0 & 0 \end{pmatrix}. \tag{19}$$

Note that here we have only assumed A_{ff} to be nonsingular. An interesting result of (19) is that the error after $\ell > 0$ iterations only depends on previous error at F-points, that is, at each iteration, C-point error is eliminated and then acquired based on error at F-points. In the case of lower triangular A_{ff}, (18) can act as a preconditioner with very similar properties to AMGir. Although we do not numerically explore preconditioning with $\widehat{M}$ here, its theoretical similarity to AMGir is worth demonstrating for future work.

Remark 8. An infinite Neumann expansion (11) can be used to approximate A_{ff}^{-1} for general non-triangular matrices, and similar algorithms developed (16) and (19). However, in the case of triangular matrices, an order-k Neumann approximation to A_{ff}^{-1} is exact on the diagonal and first k sub-diagonals. For general matrices, such accuracy cannot be guaranteed. Even in the case of symmetric positive definite matrices, where off-diagonal entries decay exponentially fast for well-conditioned A_{ff} [12], a truncated Neumann expansion does not necessarily provide a good approximation to A_{ff}^{-1} (or its diagonal for that matter). The application of AMGir to matrices with non-triangular components is on-going work.

4. Algorithm. Algorithmically, AMGir takes on the traditional structure of an AMG method. Building on the theoretical motivation provided in Section 3, here we develop a practical AMGir algorithm, with particular attention paid to limiting the setup and solve complexity, while retaining strong convergence.

4.1. Blocks. In certain cases, matrix equations solved by AMG admit a natural block structure, the most common of which are from systems of PDEs, where a single spatial node has several variables, say k, discretized over it. Two traditional approaches to handling such a system are (i) to treat each variable as a block, resulting in a decomposition of the matrix A into $n/k \times n/k$ blocks, or (ii) to treat each node over which multiple variables are discretized as a block or "super-node," resulting in a decomposition of the matrix A into $k \times k$ blocks. For the super-node approach, coarsening, interpolation, etc., are all done in a block sense, e.g., each block of k DOFs is designated either a C-block or an F-block [48]. A conceptually similar approach was used for discontinuous discretizations of elliptic PDEs in [41]. Typically, adjacent elements in a discontinuous discretization each have an independent DOF defined on the same spatial node and, thus, [41] defines a block in the initial matrix as all nodes defined at a given spatial location.

The problems considered here are not systems and do not contain super-nodes in the traditional sense, but discontinuous discretizations do lead to an inherent block structure. Recall that AMGir is designed for matrices with a triangular structure. Discontinuous upwind discretizations result in a block lower-triangular matrix in some ordering, where each block corresponds to a discontinuous finite element. Three approaches for AMGir applied to $A\mathbf{x} = \mathbf{b}$ are then as follows:

(a) Ignore block structure and treat $A\mathbf{x} = \mathbf{b}$ as a scalar problem.
(b) Let D_b be the block diagonal of A, and solve $D_b^{-1}A\mathbf{x} = D_b^{-1}\mathbf{b}$, that is, scale the block diagonal of A to consist of identity blocks.
(c) Treat $A\mathbf{x} = \mathbf{b}$ as a block system, that is, perform coarsening, relaxation, interpolation, and restriction in a block setting.

Ignoring block structure altogether achieves reasonable convergence factors for some problems, but also diverges for others. As it turns out, some consideration of the element-wise block structure is generally necessary for strong convergence of AMGir. Such results highlight the significance of A being triangular for successful reduction. This leaves options (b) and (c) for handling the block structure. With respect to convergence, all tests have indicated comparable convergence rates for these two options. However, setup and cycle complexities increase when using the full block-approach (c) as opposed to scaling by the block-inverse and solving as a scalar problem (b). For this reason, the first step of AMGir scales by the block-inverse of A, after which everything is done in a scalar setting. It is worth pointing out that in terms of finite-element order, the cost of constructing high-order discretizations becomes intractable far before the cost of computing D_b^{-1}, that is, computing D_b^{-1} is not a limiting factor for finite element order.

Remark 9. Depending on the discretization, scaling by D_b^{-1} can increase the number of nonzeros in the matrix by up to 1.5 times, thus increasing complexities. This is accounted for in numerical results presented in Section 5, that is, presented complexities are based on the number of nonzeros in the *initial matrix*.

4.2. Matrix truncation. For scalable convergence of AMGir in serial and in parallel, it is important to limit the complexity of coarse-grid operators. This motivates building R only based on strong connections as discussed in Section 4.5, and

also motivates truncating or lumping of operators as discussed here. The idea is very simple: remove entries from a matrix in the hierarchy, A_ℓ, that are smaller than some threshold, typically with respect to the diagonal element of the given row. Such methods have been used in AMG for symmetric problems with diffusive components (see, for example, [9, 21, 54]). Heuristically, eliminating small entries is even more appropriate in the hyperbolic setting, because the solution at any given point only depends on the solution at other points upwind along the characteristic. In the discrete setting, small entries that arise in matrix operations are often not aligned with the characteristic and are more of a numerical effect, suggesting that some can be eliminated without a significant degradation of convergence. In AMGir, elements $\{a_{ij} \mid j \neq i, |a_{ij}| \leq \varphi |a_{ii}|\}$ are eliminated (that is, set to zero) for each row i of matrix A_ℓ, and some drop-tolerance φ.

Remark 10. In a more traditional AMG approach, one can also eliminate small off-diagonal entries from a matrix by "lumping" or adding them to the diagonal. The purpose of lumping is to preserve the row-sum of the matrix; however, numerical tests indicate no improvement in performance by lumping entires over eliminating them, and thus it is not explored further in this work.

4.3. Coarsening. Recall that the staple of AMGir is approximating $-A_{cf}A_{ff}^{-1}$ in ideal restriction and coupling coarse-grid correction with relaxation on F-points. The effectiveness of each of these processes depends heavily on A_{ff} being well-conditioned and diagonally dominant. Such goals are consistent with the goals of coarsening in classical AMG [10]. Thus, the so-called classical strength of connection (SOC) and classical AMG CF-splitting are used as the coarsening procedure for AMGir [10, 48]. For coarsening, the SOC takes the hard minimum approach, that is, the set of strong connections for a given row consists of negative entries that are large with respect to the largest negative off-diagonal in that row:

$$\mathcal{N}_\ell = \{a_{\ell k} \mid -a_{\ell k} \geq \psi \max_{j \neq \ell} -a_{\ell j}\}.$$

Various alternative SOC measures and coarsening routines were tested as well, such as an absolute value as opposed to a hard minimum, a symmetric smoothed aggregation measure, and an evolution measure [40], but the classical measure emphasizing large negative values consistently performs best. Once a SOC matrix has been built, this is used to compute a classical AMG CF-splitting of DOFs [10], and the algorithm proceeds to approximating R_{ideal}.

4.4. Relaxation. Because ideal restriction is designed to eliminate error at C-points (Theorem 2), coarse-grid correction was coupled with Jacobi F-relaxation (17) in designing and analyzing AMGir (Section 3). One F-Jacobi iteration takes the form

$$\mathbf{x}_f^{i+1} = \mathbf{x}_f^i + D_{ff}^{-1}(\mathbf{b}_f - A_{ff}\mathbf{x}_f^i - A_{fc}\mathbf{x}_c^i), \tag{20}$$

where D_{ff} is the diagonal of A_{ff}. Interestingly, error propagation for ℓ F-Jacobi iterations (17) exactly takes the form of using an ℓth-order Neumann approximation to interpolate a coarse-grid correction. In both cases, an accurate coarse-grid correction at C-points is used to reduce error at F-points. In fact, given a correction term $(RAP)^{-1}RA\mathbf{e}$ to be interpolated to the fine grid, the following two algorithms are equivalent:

- Interpolate the correction using $\hat{P} = \begin{pmatrix} -\Delta^{(\ell)} A_{fc} \\ I \end{pmatrix}$.

- Interpolate the correction to C-points only using $\hat{P} = \begin{pmatrix} 0 \\ I \end{pmatrix}$, but perform ℓ iterations of F-Jacobi relaxation on the correction before updating the solution.

Stronger still in terms of error reduction is performing ℓ iterations of F-Jacobi relaxation on the system, following the correction. Given this equivalence, AMGir relies on F-relaxation to reduce error at F-points, and uses a very sparse and low-order approximation to P_{ideal} to limit OC and CC. This is discussed in Section 4.6.

4.5. Restriction. As mentioned in Section 4.2, controlling the sparsity-pattern fill-in of coarse-grid operators is important in controlling complexity of the method, and this fill-in depends on the sparsity of R and P. Thus, in practice, restriction operators are built in AMGir by using a truncated Neumann expansion (11) on a SOC matrix, as opposed to directly on A, to reduce the number of nonzeros in R. Here, strong connections are defined through a classical SOC based on the absolute value, that is, strong connections to node i are given by the set

$$\widehat{\mathcal{N}}_i = \{a_{ij} \mid |a_{ij}| \geq \phi \max_{\ell \neq j} |a_{\ell j}|\}. \tag{21}$$

Restriction is then built as follows: let C be the SOC matrix associated with (21). Then $A_{cf}A_{ff}^{-1}$ is approximated by $C_{cf}\widehat{\Delta}^{(k)}$, where $\widehat{\Delta}^{(k)}$ is a kth-order Neumann approximation to C_{ff}^{-1} (11).

4.6. Interpolation. As mentioned in Section 4.4, AMGir can and does rely primarily on relaxation to reduce error at F-points. Thus, the main role of interpolation is in defining the coarse-grid operator, $\mathcal{K} := RAP$. Continuing with the goal of limiting coarse-grid fill in, we want to keep P as sparse as possible, while retaining certain properties important to AMGir. Specifically, we want to maintain the constant vector as an algebraically smooth mode in $\mathcal{K}$, because this is a fundamental assumption to the chosen SOC and coarsening routine (Section 4.3). This motivates an *injection-based* interpolation. Injection interpolation determines each F-point value from, and only from, the value at its strongest-connected C-point neighbor (like a truncated A_{fc}). In this case, each row of P has exactly one nonzero, equal to one.

Remark 11. Although using a more accurate approximation to P_{ideal} for P also gives a more accurate coarse-grid operator with respect to the Schur complement, in practice the improvement in convergence of such an approach does not overcome the increase in CC.

4.7. Parameter choice. More robust AMG solvers such as root-node AMG [33], which are designed to solve a wider class of problems than traditional AMG methods, have many parameters to tune for optimal convergence. One of the novel features of AMGir is its relative invariance to parameter choice, that is, a "good" set of parameters has proved effective on all discretizations and problems that AMGir has been applied to. To demonstrate this robustness, all results are presented with a fixed set of parameter values. The matrix A is initially scaled by a block-diagonal inverse, after which the system is treated as a scalar problem. No Krylov acceleration is used in the solve phase, and $V(0, 2_F)$-cycles are applied, with two iterations of Jacobi F-relaxation applied as post-relaxation. Classical AMG coarsening is used to generate a CF-splitting, based on a classical SOC with threshold $\psi = 0.3$. The truncation tolerance is chosen as $\varphi = 10^{-3}$, which is numerically motivated in Section 5.3. A degree-one Neumann expansion (11) is used to build R with SOC tolerance $\phi = 0.025$, and so-called injection interpolation is used (Section 4.6).

5. Numerical results.

5.1. Test problems and discretizations. The model problem used here is the steady-state transport equation:

$$\begin{aligned} \mathbf{b}(x,y)\cdot\nabla u + c(x,y)u &= q(x,y) \quad \Omega, \\ u &= g(x,y) \quad \Gamma_{\text{in}}, \end{aligned} \tag{22}$$

for domain Ω and inflow boundary Γ_{in}. Multiple cases are studied that encompass spatially dependent source terms, $q(x,y)$, discontinuities in the material coefficient, $c(x,y)$, and constant and non-constant flow direction, $\mathbf{b}(x,y)$, over structured and unstructured meshes. Two domains and the respective solution for a constant flow are given in Figure 2. Several variations of the *inset* domain with non-constant flow are shown in Figure 4 in Section 5.2.

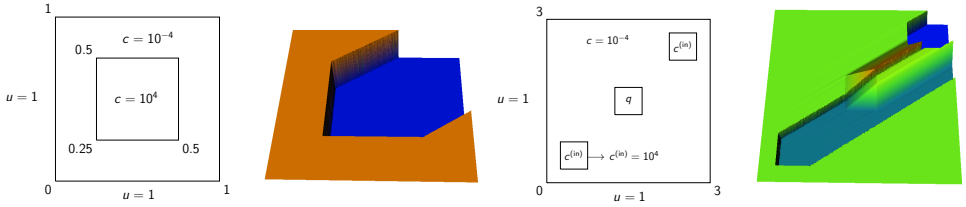

(a) *Inset* domain (b) *Block-source* domain

Fig. 2: Two domains for the steady-state transport equation. Inflow boundaries consist of the south and west boundaries with inflow $u = 1$. Material coefficient $c(x,y)$ is piecewise constant in both cases, with changes of eight orders of magnitude. The *block-source* domain has an interior source $q(x,y) = 1$ in the interior block and $q(x,y) = 0.5 \cdot 10^4$ in the lower left block.

To accompany the different domains considered, multiple upwind discretizations are implemented. A first-order lumped finite element discretization [35, 36] is applied on structured and unstructured meshes. Standard fully upwinded discontinuous Galerkin (DG) discretizations [27, 46] are also tested, with finite element orders $0-6$, and the boundary conditions strongly- and weakly-imposed. A comprehensive introduction can be found in [17]. Standard upwinded DG methods arise as special cases in [14] and for almost-scattering-free problems in [45]. In all problems tested, AMGir performs near identically on strongly and weakly enforced boundary conditions; thus, moving forward, weakly-enforced boundary conditions are (arbitrarily) used.

5.2. Non-constant flow and angular variation. Anisotropies on unstructured meshes and strong non-grid-aligned anisotropies can prove difficult for AMG solvers [33, 51]. An additional novel feature of AMGir is its robustness with respect to problems defined on structured and unstructured meshes. Figure 3 demonstrates this robustness for LCB discretizations of the *inset* problem on structured and unstructured meshes, with fixed angle $\Omega := \mathbf{b}(x,y) = (\cos(\theta), \sin(\theta))$, for angles $\theta \in [0, \pi/2]$.

Because unstructured meshes are often used in practice and typically more difficult from a solver perspective, further results in Section 5 are presented for unstructured meshes. By nature of the relative invariance of convergence and complexity with respect to angle, further tests in Section 5 arbitrarily fix $\theta = 3\pi/16$ for the *inset* and block-source problems.

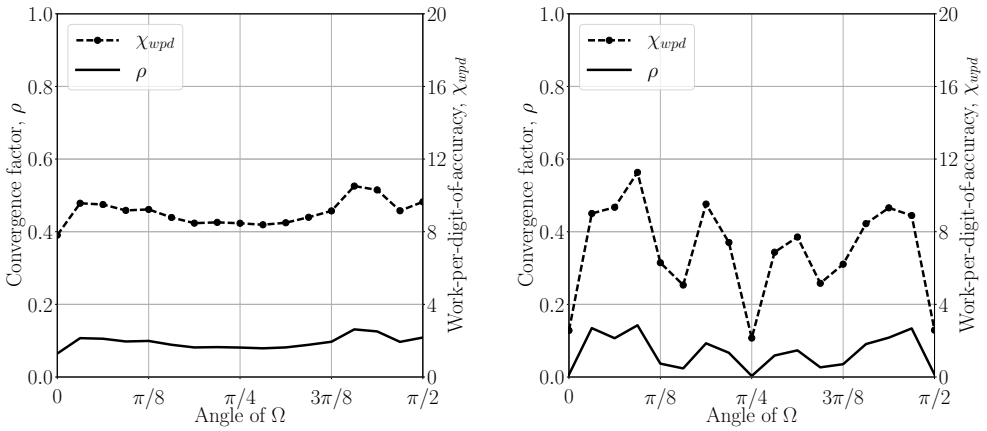

(a) LCB unstructured mesh

(b) LCB structured mesh

Fig. 3: Convergence factors and work-per-digit-of-accuracy for AMGir applied to LCB discretizations of the *inset* problem, with angles between 0 and $\pi/2$, on unstructured and structured meshes, and ≈ 2.25M DOFs.

Remark 12. For some angles, AMGir performs better in terms of convergence factor and work per digit of accuracy on an unstructured mesh compared with a structured mesh. However, even when these measures favor an unstructured mesh, the wall-clock time of the setup and solve phase is at least $2\times$ faster in *all cases* for a structured mesh over an unstructured mesh. It is possible that a structured mesh makes for a more structured matrix amenable to matrix-vector operations, but a detailed analysis is outside the scope of this work.

In addition to being robust with respect to angular variations, AMGir is generally insensitive to flow direction $\mathbf{b}(x, y)$. Figure 4 shows the solution to three different non-constant $\mathbf{b}(x, y)$ defined on the *inset* domain. Table 1 then shows the convergence factor, CC, and work-per-digit-of-accuracy of AMGir applied to a constant flow, and each of the moving flows shown in Figure 4.

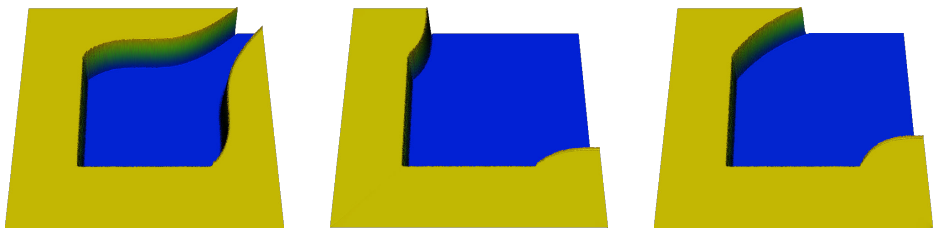

(a) $\mathbf{b}_1(x, y) = (\cos(\pi y)^2, \cos(\pi x)^2)$.

(b) $\mathbf{b}_2(x, y) = (\sin(\pi y)^2, \sin(\pi x)^2)$.

(c) $\mathbf{b}_3(x, y) = (y^4, \cos(\pi x/2)^2)$.

Fig. 4: Solution of three different moving flows defined on the *inset* domain.

$\mathbf{b}(x,y)$	Ω	$\mathbf{b}_1(x,y)$	$\mathbf{b}_2(x,y)$	$\mathbf{b}_3(x,y)$
ρ	0.10	0.11	0.09	0.10
CC	7.58	7.49	7.54	7.51
χ_{WPD}	7.56	7.86	7.17	7.63

Table 1: Convergence factor, CC, and work-per-digit-of-accuracy for AMGir applied to variations in flow direction, $\mathbf{b}(x,y)$, and constant flow direction $\mathbf{b}(x,y) = \Omega = (\cos(3\pi/16), \sin(3\pi/16))$. Discretizations are defined on the *inset* domain with an unstructured mesh, using upwind DG with linear elements and ≈ 2.7M DOFs.

5.3. Truncation and building restriction. Truncating coarse-grid operators as introduced in Section 4.2 offers a significant improvement in AMGir performance when considering work-per-digit-of-accuracy. One of the difficulties with scaling AMGir to progressively larger problems is the fill-in of the coarse-grid operators, which can greatly increase the CC. However, as discussed in Section 4.2, many of the non-zero entries in A_i for i lower in the AMGir hierarchy are likely to be relatively small with respect to diagonal elements, and can be removed without substantial degradation in convergence. Figure 5 shows work-per-digit-of-accuracy as a function of truncation tolerance for many combinations of problems considered here. All discretizations, all domains, and finite elements of order $0-6$ are tested and shown in Figure 5. Despite matrices with very different element magnitude and connectivity, a truncation tolerance of $\varphi = 10^{-3}$ (shown in dotted black) is an effective choice in all cases, leading to a $20-50\%$ reduction in iteration cost. Setup cost is not measured in detail, but setup wall-clock times are reduced by comparable amounts.

Remark 13. It is worth pointing out that $\varphi = 10^{-3}$ is not the *best* choice of truncation tolerance for all problems considered, only that it is an effective choice in all cases. For matrices with high connectivity such as higher order finite elements, a larger φ reduces the work-per-digit-of-accuracy. Conversely, matrices with very low

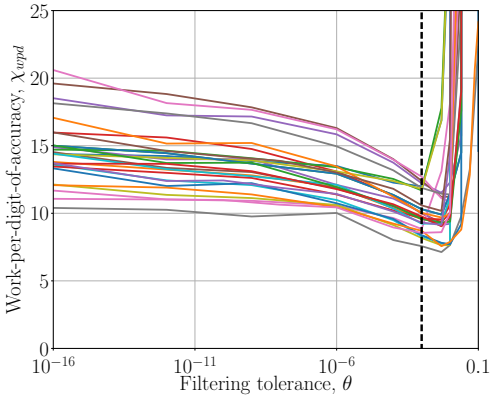

Fig. 5: Work-per-digit of accuracy as a function of truncation tolerance for AMGir applied to various discretizations of the steady-state transport equation. The dotted block line shows $\varphi = 10^{-3}$, which is an effective choice for all problems tested.

connectivity, such as a degree zero DG discretization, favor smaller φ.

5.4. Scaling in h and element order. The final numerical study presented here is the scaling of AMGir with respect to DOFs, as well as finite-element degree. One of the exciting features of AMGir is its ability to solve high-order finite-element discretizations, something that AMG methods often struggle with. This is done in two parts: first, a finite-element degree is fixed and the performance of AMGir is examined as spatial resolution tends to zero; and, second, the spatial resolution is fixed and AMGir is applied to finite element discretizations of degree $d = 1, ..., 6$. Two cycle types will be considered: a V-cycle and an F-cycle. An F-cycle consists of first restricting all the way to the coarsest grid (as in a V-cycle) without relaxing along the way, followed by performing an additional V-cycle at each level of the hierarchy when traversing back to the finest grid. The complexity of an F-cycle is typically 1.5–3 times that of a V-cycle. Although F-cycles originate in full multigrid and full AMG, which focus on achieving discretization-level accuracy in a single F-cycle, accuracy with respect to discretization is not considered in this work. Instead, the F-cycle is used as it can provide more robust convergence and scaling than a V-cycle, at a much lower cost than alternatives such as W-cycles and K-cycles.

Figure 6 shows AMGir scaling of work-per-digit-of-accuracy and convergence fac-

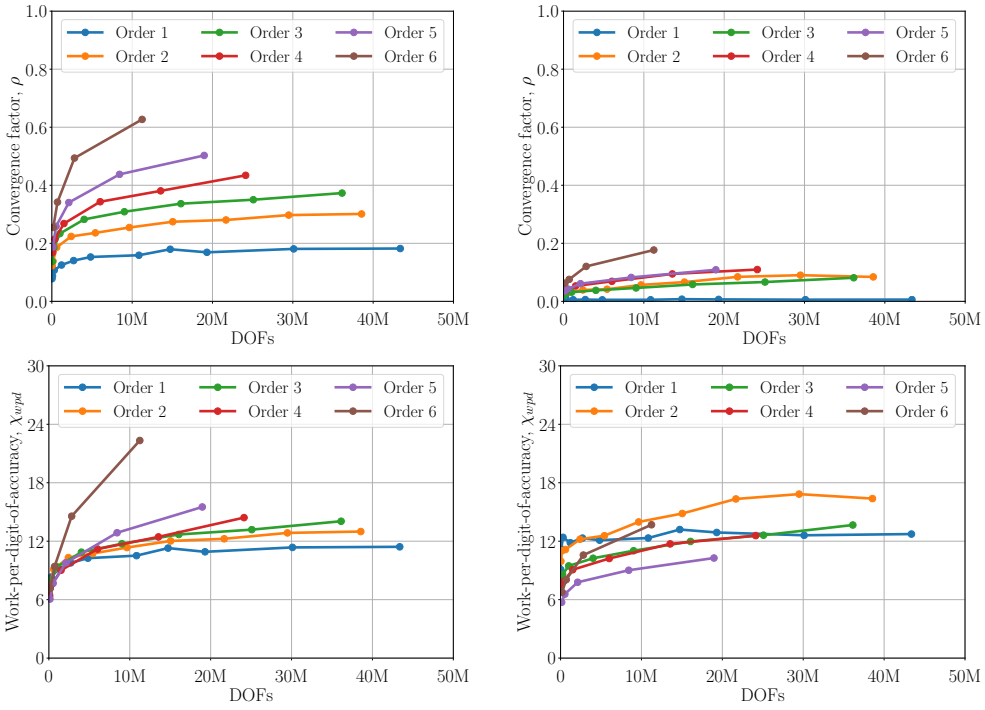

(a) V-cycle (b) F-cycle

Fig. 6: Scaling of convergence factor and work-per-digit-of-accuracy as a function of DOFs, for AMGir applied to upwind DG discretizations of the block-source problem, with finite-element degrees 1–6.

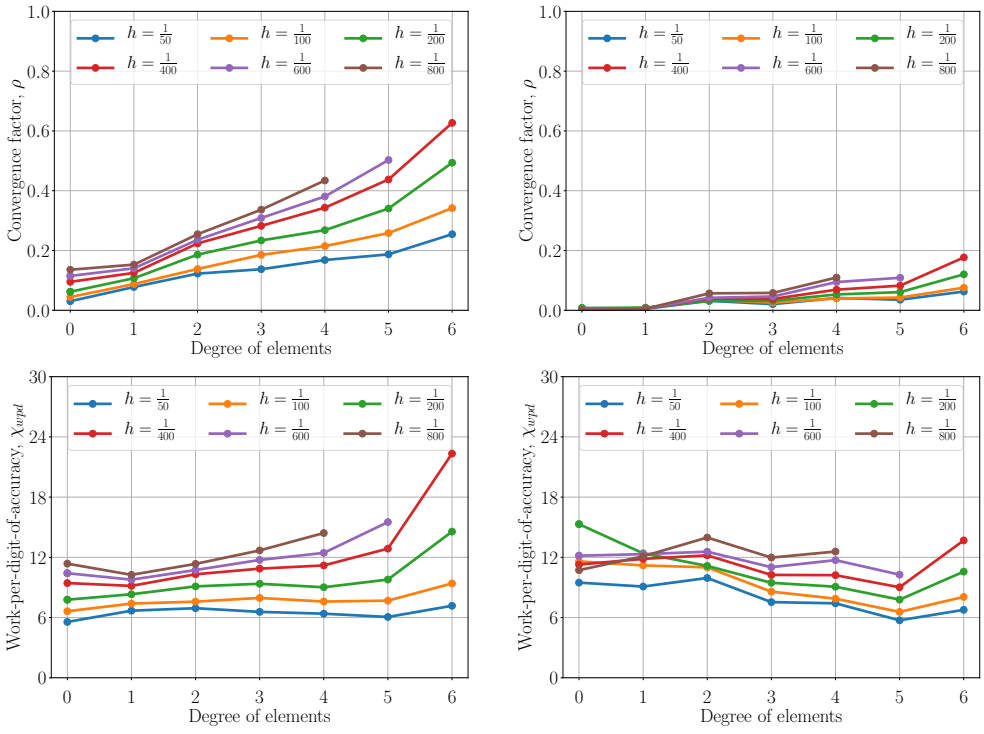

(a) V-cycle (b) F-cycle

Fig. 7: Scaling of convergence factor and work-per-digit-of-accuracy as a function of finite-element degree, for AMGir applied to upwind DG discretizations of the block-source problem, with spatial resolution $h = \frac{1}{50}$ to $h = \frac{1}{800}$.

tor as a function of number of DOFs. Low-order elements have reached asymptotic behavior in V- and F-cycles, thus showing AMGir to be a scalable method up to $\approx$ 40M DOFs. Convergence factors and work-per-digit-of-accuracy are still slowly increasing for higher-order elements; however, memory constraints in the serial setting prevented scaling to larger problems. Figure 7 shows scaling of convergence and work-per-digit-of-accuracy as a function of finite element order. Notice that for F-cycles in particular, scaling with respect to work is near-perfect for all spatial resolutions tested. Based on Figure 6, it is unlikely for such perfect scaling to continue as $h \to 0$; nevertheless, it is encouraging that AMGir performs so well on high-order elements.

6. Conclusions. This work develops a reduction-based AMG method for matrices with triangular structure based on the concept of ideal restriction. Theory is developed showing that, in the spectral sense, AMGir is a robust preconditioner for matrices with triangular structure, and is guaranteed to converge to the exact solution in less than n iterations. Numerical tests then show AMGir to actually be a fast solver of upwind discretizations, which does not require the parameter tuning that some AMG methods require. For a single set of parameters, AMGir is able to achieve strong convergence factors on multiple domains and high-order upwind discretizations of the steady-state transport equation.

The next step in studying AMGir is a parallel implementation to study of its parallel efficiency. This will also allow further testing of asymptotic behavior for high-order elements. So far, this work has been confined to matrices with block-triangular structure for reasonably small blocks. Future directions involve studying AMGir applied to other block-triangular matrices, block-triangular matrices with blocks too large to explicitly invert, or matrices with some symmetric or non-triangular component.

REFERENCES

[1] M. L. Adams and E. W. Larsen, *Fast iterative methods for discrete-ordinates particle transport calculations*, Progress in nuclear energy, 40 (2002), pp. 3–159.

[2] M. P. Adams, M. L. Adams, and W. D. Hawkins, *Provably optimal parallel transport sweeps on regular grids*, in International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering, Idaho, 2013, Texas A and M University, College Station, United States, pp. 2535–2553.

[3] F. L. Alvarado and R. Schreiber, *Optimal parallel solution of sparse triangular systems*, SIAM Journal on Scientific Computing, 14 (1993), pp. 446–460.

[4] S. Amarala and J. W. L. Wan, *Multigrid Methods for Systems of Hyperbolic Conservation Laws*, Multiscale Modeling & Simulation, 11 (2013), pp. 586–614.

[5] E. Anderson and Y. Saad, *Solving sparse triangular linear systems on parallel computers*, International Journal of High Speed Computing, 1 (1989), pp. 73–95.

[6] T. S. Bailey and R. D. Falgout, *Analysis of massively parallel discrete-ordinates transport sweep algorithms with collisions*, in International Conference on Mathematics, Computational Methods, and Reactor Physics, Saratoga Springs, NY, 2009, Lawrence Livermore National Laboratory, Livermore, United States, pp. 1751–1765.

[7] A. H. Baker, R. D. Falgout, T. V. Kolev, and U. M. Yang, *Scaling Hypre's Multigrid Solvers to 100,000 Cores*, in High-Performance Scientific Computing, Springer London, London, 2012, pp. 261–279.

[8] W. N. Bell, L. N. Olson, and J. B. Schroder, *PyAMG: Algebraic multigrid solvers in Python v3.0*, 2015, http://www.pyamg.org. Release 3.0.

[9] A. Bienz, R. D. Falgout, W. Gropp, and L. N. Olson, *Reducing Parallel Communication in Algebraic Multigrid through Sparsification*, SIAM Journal on Scientific Computing, 38 (2016), pp. S332–S357.

[10] A. Brandt, S. F. McCormick, and J. Huge, *Algebraic multigrid (amg) for sparse matrix equations*, Sparsity and its Applications, 257 (1985).

[11] J. J. Brannick, A. Frommer, and K. Kahl, *Adaptive reduction-based multigrid for nearly singular and highly disordered physical systems*, Electronic transactions on numerical analysis, 37 (2010), pp. 276–295.

[12] J. J. Brannick and L. T. Zikatanov, *Algebraic Multigrid Methods Based on Compatible Relaxation and Energy Minimization*, in Domain Decomposition Methods in Science and Engineering XVI, Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, pp. 15–26.

[13] M. Brezina, T. A. Manteuffel, S. F. McCormick, J. W. Ruge, and G. Sanders, *Towards Adaptive Smoothed Aggregation (αSA) for Nonsymmetric Problems*, SIAM Journal on Scientific Computing, 32 (2010), pp. 14–39.

[14] F. Brezzi, L. D. Marini, and E. Süli, *Discontinuous Galerkin methods for first-order hyperbolic problems*, Mathematical models and methods in applied sciences, 14 (2004), pp. 1893–1903.

[15] L. M. Carvalho, L. Giraud, and P. Le Tallec, *Algebraic two-level preconditioners for the Schur complement method*, SIAM Journal on Scientific Computing, 22 (2001), pp. 1987–2005.

[16] G. Colomer, R. Borrell, F. X. Trias, and I. Rodríguez, *Parallel algorithms for Sn transport sweeps on unstructured meshes*, Journal of Computational Physics, 232 (2013), pp. 118–135.

[17] D. Di Pietro and A. Ern, *Mathematical Aspects of Discontinuous Galerkin Methods*, Mathématiques et Applications, Springer Berlin Heidelberg, 2011.

[18] V. Dobrev, T. Kolev, and N. A. Petersson, *Two-level convergence theory for parallel time integration with multigrid*, SIAM Journal on Scientific Computing, (2016).

[19] R. Falgout, T. Manteuffel, B. O'Neill, and J. Schroder, *Multigrid Reduction in Time for Nonlinear Parabolic Problems*, Numerical Linear Algebra with Applications, ((in review)).

[20] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, and J. B. Schroder, *Par-*

allel Time Integration with Multigrid, SIAM Journal on Scientific Computing, 36 (2014), pp. C635–C661.

[21] R. D. Falgout and J. B. Schroder, *Non-Galerkin coarse grids for algebraic multigrid*, SIAM Journal on Scientific Computing, 36 (2014), pp. C309–C334.

[22] R. D. Falgout and P. S. Vassilevski, *On Generalizing the Algebraic Multigrid Framework*, SIAM Journal on Numerical Analysis, 42 (2004), pp. 1669–1693.

[23] R. D. Falgout and U. M. Yang, *hypre: A library of high performance preconditioners*, European Conference on Parallel Processing, 2331 LNCS (2002), pp. 632–641.

[24] H. Foerster, K. Stüben, and U. Trottenberg, *Non-standard multigrid techniques using checkered relaxation and intermediate grids*, Academic Press, New York, (1981), pp. 285–300.

[25] M. J. Gander and S. Vandewalle, *Analysis of the Parareal Time-Parallel Time-Integration Method*, SIAM Journal on Scientific Computing, 29 (2007), pp. 556–578.

[26] W. Hawkins, *Efficient massively parallel transport sweeps*, in Transactions of the American Nuclear Society, Texas A and M University, College Station, United States, Dec. 2012, pp. 477–481.

[27] P. Lesaint and P. Raviart, *On a finite element method for solving the neutron transport equation*, C. de Boor (Ed.), Mathematical Aspects of Finite Elements in Partial Differential Equations, (1974), pp. 89–123.

[28] R. Li and Y. Saad, *GPU-accelerated preconditioned iterative linear solvers*, The Journal of Supercomputing, 63 (2013), pp. 443–466.

[29] W. Liu, A. Li, J. Hogg, I. S. Duff, and B. Vinter, *A Synchronization-Free Algorithm for Parallel Sparse Triangular Solves*, European Conference on Parallel Processing, 9833 (2016), pp. 617–630.

[30] S. P. MacLachlan, T. A. Manteuffel, and S. F. McCormick, *Adaptive reduction-based AMG*, Numerical Linear Algebra with Applications, 13 (2006), pp. 599–620.

[31] J. Mandel, *On block diagonal and Schur complement preconditioning*, Numerische Mathematik, 58 (1990), pp. 79–93.

[32] T. A. Manteuffel, S. Münzenmaier, and B. S. Southworth, *Scaling and Solving the Self-Adjoint Form for Steady-State Transport*, (in preparation).

[33] T. A. Manteuffel, L. N. Olson, J. B. Schroder, and B. S. Southworth, *A root-node based algebraic multigrid method*, SIAM Journal on Scientific Computing, (accepted) (2017).

[34] C. Mense and R. Nabben, *On algebraic multi-level methods for non-symmetric systems – Comparison results*, Linear Algebra and its Applications, 429 (2008), pp. 2567–2588.

[35] J. E. Morel and J. S. Warsa, *An S_n Spatial Discretization Scheme for Tetrahedral Meshes*, Nuclear science and engineering, 151 (2005), pp. 157–166.

[36] J. E. Morel and J. S. Warsa, *Spatial Finite-Element Lumping Techniques for the Quadrilateral Mesh S_n Equations in X-Y Geometry*, Nuclear science and engineering, 156 (2007), pp. 325–342.

[37] M. F. Murphy, G. H. Golub, and A. J. Wathen, *A note on preconditioning for indefinite linear systems*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1969–1972.

[38] Y. Notay, *A robust algebraic multilevel preconditioner for non-symmetric M-matrices*, Numerical Linear Algebra with Applications, 7 (2000), pp. 243–267.

[39] Y. Notay, *Algebraic analysis of two-grid methods: The nonsymmetric case*, Numerical Linear Algebra with Applications, 17 (2010), pp. 73–96.

[40] L. N. Olson and J. B. Schroder, *A new perspective on strength measures in algebraic multigrid*, Numerical Linear Algebra with Applications, 17 (2010), pp. 713–733.

[41] L. N. Olson and J. B. Schroder, *Smoothed aggregation multigrid solvers for high-order discontinuous Galerkin methods for elliptic problems*, Journal of Computational Physics, 230 (2011), pp. 6959–6976.

[42] L. N. Olson, J. B. Schroder, and R. S. Tuminaro, *A General Interpolation Strategy for Algebraic Multigrid Using Energy Minimization*, SIAM Journal on Scientific Computing, 33 (2011), pp. 966–991.

[43] C. W. Oosterlee, F. J. Gaspar, T. Washio, and R. Wienands, *Multigrid Line Smoothers for Higher Order Upwind Discretizations of Convection-Dominated Problems*, Journal of Computational Physics, 139 (1998), pp. 274–307.

[44] H. Park, D. A. Knoll, R. M. Rauenzahn, C. K. Newman, J. D. Densmore, and A. B. Wollaber, *An Efficient and Time Accurate, Moment-Based Scale-Bridging Algorithm for Thermal Radiative Transfer Problems*, SIAM Journal on Scientific Computing, 35 (2013), pp. S18–S41.

[45] J. C. Ragusa, J. L. Guermond, and G. Kanschat, *A robust SN-DG-approximation for radiation transport in optically thick and diffusive regimes*, Journal of Computational Physics,

231 (2012), pp. 1947–1962.
[46] W. Reed and T. Hill, *Triangular mesh methods for the neutron transport equation*, Tech. Rep. LA-UR-73-479, Los Alamos Scientific Laboratory, Los Alamos, NM, (1973).
[47] M. Ries, U. Trottenberg, and G. Winter, *A note on mgr methods*, Linear Algebra and its Applications, 49 (1983), pp. 1–26.
[48] J. W. W. Ruge and K. Stüben, *Algebraic multigrid*, Multigrid methods, 3 (1987), pp. 73–130.
[49] Y. Saad and M. Sosonkina, *Distributed Schur complement techniques for general sparse linear systems*, SIAM Journal on Scientific Computing, 21 (1999), pp. 1337–1356.
[50] M. Sala and R. S. Tuminaro, *A New Petrov–Galerkin Smoothed Aggregation Preconditioner for Nonsymmetric Linear Systems*, SIAM Journal on Scientific Computing, 31 (2008), pp. 143–166.
[51] J. B. Schroder, *Smoothed aggregation solvers for anisotropic diffusion*, Numerical Linear Algebra with Applications, 19 (2012), pp. 296–312.
[52] M. Sea d and A. Klar, *Efficient preconditioning of linear systems arising from the discretization of radiative transfer equation*, in Challenges in scientific computing—CISC 2002, Springer, Berlin, Berlin, Heidelberg, 2003, pp. 211–236.
[53] B. S. Southworth, J. J. Brannick, S. P. MacLachlan, J. B. Schroder, and L. T. Zikatanov, *The Role of Energy Minimization in Algebraic Multigrid Interpolation*, tech. report, Lawrence Livermore National Laboratory, in preparation.
[54] E. Treister and I. Yavneh, *Non-Galerkin Multigrid Based on Sparsified Smoothed Aggregation*, SIAM Journal on Scientific Computing, 37 (2015), pp. A30–A54.
[55] P. S. Vassilevski, *Multilevel Block Factorization Preconditioners*, Matrix-based Analysis and Algorithms for Solving Finite Element Equations, Springer Science & Business Media, Oct. 2008.
[56] T. A. Wiesner, R. S. Tuminaro, W. A. Wall, and M. W. Gee, *Multigrid transfers for nonsymmetric systems based on Schur complements and Galerkin projections*, Numerical Linear Algebra with Applications, 21 (2013), pp. 415–438.
[57] I. Yavneh, C. H. Venner, and A. Brandt, *Fast multigrid solution of the advection problem with closed characteristics*, SIAM Journal on Scientific Computing, 19 (1998), pp. 111–125.